

EXHIBIT 17

DATA COMMUNICATIONS

FACILITIES, NETWORKS, AND SYSTEMS DESIGN

Dixon R. Doll

President
DMW Telecommunications Corporation
Ann Arbor, Michigan
and
Adjunct Faculty Member
IBM Systems Research Institute
New York, New York

A Wiley-Interscience Publication
JOHN WILEY & SONS
New York • Chichester • Brisbane • Toronto

Copyright © 1978 by John Wiley & Sons, Inc.

All rights reserved. Published simultaneously in Canada.

Reproduction or translation of any part of this work beyond that permitted by Sections 107 or 108 of the 1976 United States Copyright Act without the permission of the copyright owner is unlawful. Requests for permission or further information should be addressed to the Permissions Department, John Wiley & Sons, Inc.

Library of Congress Cataloging in Publication Data:

Doll, Dixon R.

Data communications.

"A Wiley-Interscience publication."

Includes bibliographical references and index.

1. Data transmission systems. 2. Computer networks.

I. Title.

TK5105.D64 384 77-12508
ISBN 0-471-21768-9

Printed in the United States of America

10 9 8 7 6 5 4 3 2

7.3. Synchronous Time-Division Multiplexing (STDM)

311

sharing becomes possible. The capacity of the voice-grade line is shared by the FDM equipment, which creates multiple independent subchannels. Each subchannel in turn may be time shared, on either a contention or a polled basis, by multiple remote terminals.

For example, imagine that two asynchronous terminals in each of the following cities—Boston, New York, Chicago, and Kansas City—are to be tied to a central computer in Los Angeles. Figure 7.3 illustrates four possible configurations, the first of which uses individual leased lines to each remote terminal. The second approach utilizes FDM equipment with an assumed capacity of four subchannels per voice-grade line and no sharing of the subchannels. The third approach, which clearly has a still lower line cost, uses FDM equipment with an assumed capacity of four subchannels per voice-grade line and two terminals per FDM subchannel. The fourth approach has the same line costs as the third configuration. However, an FDM system with an assumed capacity of eight subchannels per voice-grade line (and no sharing of subchannels) is postulated. Note that all configurations (except the one in which sharing of the subchannels is permitted) require eight ports at the central site. Only four ports are required when each subchannel can be shared by two remote terminals.

Other popular examples of the application of frequency-division multiplexing techniques are their use in special modems to create a full-duplex channel over a two-wire circuit (discussed in Chapter 2) and to provide extra low-speed teletype-grade channels on voice-grade circuits, primarily in international applications. In such situations, the collective costs of separate voice-grade and teletype (sub-voice-grade) lines are often sufficiently large that it may be less expensive to operate one leased voice-grade line (for either voice communications or data transmission up to 9600 bits/sec) concurrently with one or more independent low speed side channels on the same physical line. The analog data modems can usually be switched out in favor of a telephone at each end, enabling either voice or data transmission to take place independently of slow speed subchannel activity.

7.3. SYNCHRONOUS TIME-DIVISION MULTIPLEXING (STDM)

Time-division multiplexing devices that create a permanently dedicated time slot or subchannel for each port in the sharing group are classified as STDMs. By contrast, statistical or asynchronous TDMs dynamically allocate the subchannels or time slots on a statistical basis to increase line



efficiency by providing time slots only for ports actively transmitting data.³

As shown in Figure 7.4, STDMS share a synchronous communication line by cyclically scanning incoming data from input ports, peeling off bits or characters, and interleaving them into frames on a single high speed data stream. This effect is similar to that of a high speed conveyor belt picking up objects arriving at a common point from several lower speed belts. In utilizing a given channel, STDMS is generally more efficient than FDM since it is capable of using the entire bandwidth available.

For example, STDMS can generally operate over dedicated voice-grade lines at speeds of 4800, 7200, and 9600 bits/sec, whereas FDMs' practical limit on the same line is probably in the 2000 bits/sec speed range. Generally speaking, the multiplexed data stream is transmitted serially, bit-by-bit, at a rate governed by the circuit-signal converter combination. The split-stream modems introduced in Chapter 2 are popular devices that combine the modem and STDMS functions into the modem device when the input port speeds are integer multiples of 2400 bits/sec and the total input rate does not exceed 9600 bits/sec.

Although voice-grade lines are shared in the majority of current STDMS applications, a recent regulatory development now enables end users to multiplex broadband carrier links with customer-provided STDMS equipment. Whereas STDMS can multiplex traffic from either asynchronous (start/stop) terminals, other synchronous devices, or combinations thereof, FDMs are generally used to multiplex only asynchronous terminals, although this is not an intrinsic limitation.

Contemporary STDMS may perform either bit or character interleaving on the shared line when serving start/stop terminals exclusively. In these applications, character interleaving is usually more efficient since a modest amount of bandwidth compression is possible. The start and stop bits of each character entering the STDMS may be stripped off before the character's insertion into the frame of multiplexed data. (Figure 7.5 illustrates the technique of character-interleaved STDMS, including data characters and the encodings of various end-to-end control signals.) Any bits stripped from incoming data characters are reinserted by the demultiplexing unit before distribution of the characters to their respective

³The reader should not confuse asynchronous time-division multiplexers, in the sense of their current interpretation, with STDMS that multiplex asynchronous (start/stop) terminal devices. Throughout this chapter, "STDMS" describes the time-division technique in which dedicated subchannels are created for start/stop devices, synchronous devices, or a combination thereof. "Statistical" or "asynchronous TDM," by contrast, describes all schemes where the multiplexer creates subchannels dynamically, regardless of the type of device being multiplexed.

7.3. Synchronous Time-Division Multiplexing (STDM)

313

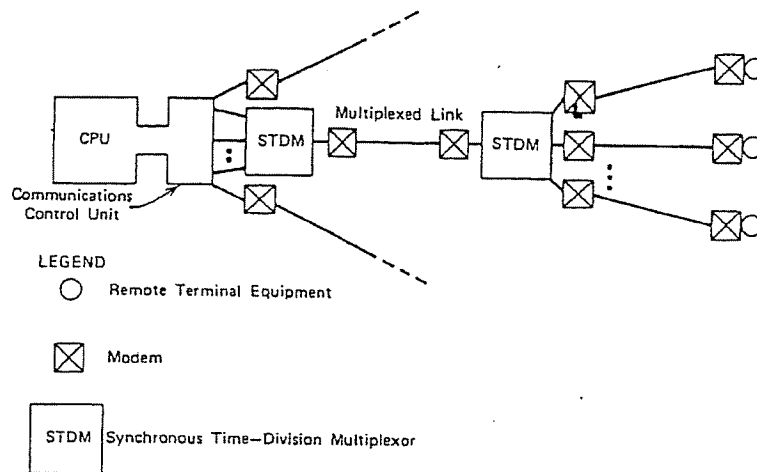


Figure 7.4. Synchronous TDMs in a teleprocessing network.

output ports. Thus the incoming characters are effectively encoded using fewer bits per character for transmission over the shared link, thereby enabling an aggregate low speed bit rate of 1.1 times the shared link's transmission rate to be accommodated in typical situations.

In newer applications involving the use of STDMS to multiplex synchronous data streams, the STDMS generally employ bit interleaving,

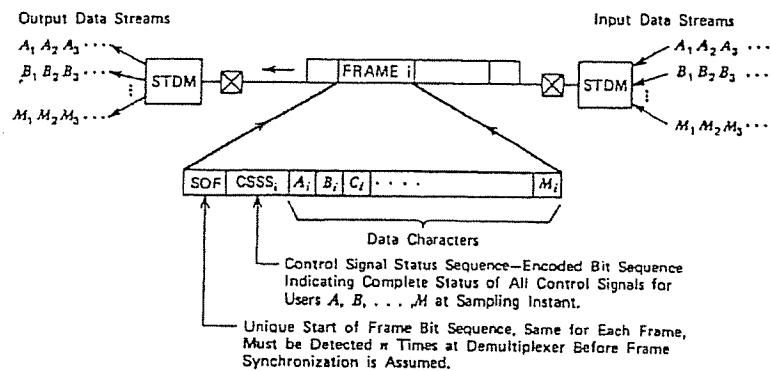


Figure 7.5. Character-Interleaved STDM and frame format.

disregarding the textual content of the incoming data streams. This data transparency is an important requirement for multiplexers being incorporated in the long haul trunks of synchronous digital data networks now being implemented by certain users and common carriers. Whether bit or character interleaving is used, special predetermined code sequences are utilized between STDMs to define the beginning of each new frame of multiplexed data. Demultiplexing is thus accomplished by the assumption of an implicit relationship between the output line or buffer address and the relative position of the time slots in an arriving frame.

When all the multiplexed terminal devices are unbuffered, the problem of fixing the scanning rates within the STDM is straightforward—the subchannel scan rates are matched to the transmission rates of the respective input lines. Ordinarily this speed corresponds, in turn, to the operating speed of the remote terminal device being served. However, when messages and message segments are queued in remote terminal buffers, the solution to the scan rate assignment problem is not so obvious. In References [3] and [4], Doll has developed a queuing theoretic design technique for determining scan rates within the STDM so that the average queuing delay experienced at the remote terminals in the sharing group is minimized.

Noise disturbances on the shared channel can cause a variety of errors, depending on whether character or bit interleaving is used. With character interleaving, an individual data bit error will at worst cause a single output character to be received in error. With bit-interleaved STDM, a similar anomaly could cause the demultiplexing unit to deliver the outputs to the wrong addresses, unless the STDMs contain their own error control capability. As a consequence, character STDMs are less sensitive than bit multiplexers to channel disturbances, although resynchronization (reestablishing the start of a data frame) takes somewhat longer than with bit-interleaved STDMs.

Higher quality STDM devices have been designed on the philosophy that random and burst errors can be allowed to cause data errors, but must virtually never cause errors in the end-to-end control signals between user terminals or in the internal network control signals between STDMs themselves. This goal may be accomplished without substantial degradation of shared link capacity by using highly redundant encodings of all vital control signals. For example, frame synchronization is obviously critical and must be preserved in the presence of noise bursts. Elaborate time averaging and/or thresholding schemes have been devised whereby frame synchronization is assumed at the demultiplexer only after a unique bit sequence has been detected a specified number of times in a given time period. Similarly, frame synchronization is assumed to be lost

Issues for Line Sharing

streams. This data is being incorporated into networks now. Whether bit or frame sequences are used, the new frame of data is the assumption of the address and

the problem is forwarded—the of the respective to the operation. However, when the buffers, the obvious. In the design the average group is

of errors, character output similar to the control capacity than bit (reestablishment) than with

topology, but the relationship between potential loadings is obvious. The network is designed after the fact in a

7.3. Synchronous Time-Division Multiplexing (STDM)

315

only when this same condition cannot be detected. Typical noise disturbances may thus be smoothed over without catastrophic effect, even though some data bit errors may occur during intervals when frame synchronization is being reestablished.

In comparison to FDMs, STDMs are expensive to cascade because a relatively complete STDM system must be used at any point where one or more subchannels are being inserted or removed. Also, when STDMs are used in cascade, the problem of coordinating the timing across multiple synchronous links must be addressed. Fundamentally, two choices exist here—to use one master clock from which all system elements derive their timing, or to use independent synchronous clocks on each link. The former alternative may be illustrated by the simple configuration shown in Figure 7.6. The modem at *A* provides the master clock signal, and the lefthand modem at *B* derives its clock from the incoming data on link *AB*. The righthand modem at *B* is slaved to its lefthand counterpart at *B* and, in turn, provides a master clock to link *BC*. With the alternative approach of independent synchronous clocks, some type of elastic buffer must be provided at node *B* to absorb data buildups caused by slight variations in the rates of the two independent clocks. Most STDM networks to date have been implemented using the single-master-clock scheme for various economic and reliability reasons.

As with FDM, an STDM configuration provides each port in the sharing group with its own dedicated appearance at the communications control unit. (See Figure 7.4.) Here it is assumed that STDMs are used in pairs, one for multiplexing and the other for demultiplexing. It is possible to eliminate the central site STDM if the communications control unit can perform the STDM function in hardware or software. (An early IBM Corporation STDM product combination known as the 2712 Multiplexor operated with a hard-wired transmission control known as the 2702, eliminating the need for two separate central site devices.)

Since most computer vendors have not emphasized time-division mul-




Figure 7.6. A synchronous TDM network.

time-division multiplexing equipment (and vice versa), users continue frequently to use the approach of paired STDMS for multiplexing and demultiplexing. It creates a well-defined hardware interface point between the communications controller and the network. It also enables conventional, vendor-supported teleprocessing control software to be utilized without modification or extension. The advantages arising from these two characteristics often outweigh the extra costs associated with the two-separate-box approach to time-division multiplexing.

Subsequent discussions of statistical multiplexing, packet switching, and intelligent communications networks based on minicomputers will reveal certain advantages of dynamic bandwidth sharing not available with conventional STDM. When such benefits can be coupled into networks with minimal requirements for modifications to the line control software, the user stands to achieve the best of both worlds—a flexible network and the full support of vendors supplying the teleprocessing control software.

7.3.1. Configuration Options in Cascaded STDM Networks

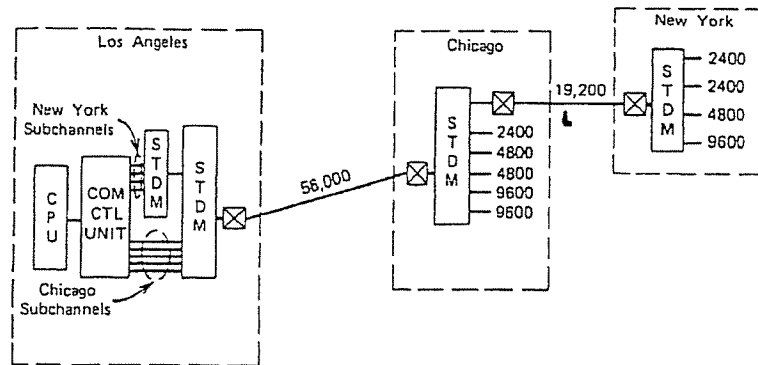
Consider a user with a Los Angeles CPU and numerous remote terminals in Chicago and New York. For illustration purposes, assume that the following terminal-speed combinations are required:

New York	Chicago
Two 2400 bits/sec terminals	One 2400 bits/sec terminal
One 4800 bits/sec terminal	Two 4800 bits/sec terminals
One 9600 bits/sec terminal	Two 9600 bits/sec terminals


Also assume that 19,200 and 56,000 bits/sec line speeds are available for use on the multiplexed STDM links; the problem is to connect all nine remote terminals into separate ports on the Los Angeles communications control unit. Figure 7.7 illustrates that the New York channels may be demultiplexed in either Chicago or Los Angeles. If demultiplexed in Chicago, the New York channels must be remultiplexed with the Chicago traffic onto the Chicago-Los Angeles link. This so-called drop-and-insert arrangement would provide convenient access to the New York channels in the Chicago site if such an arrangement is desirable. Also, extra STDM capacity could be used between New York and Chicago, if required. On the other hand, the Los Angeles demultiplex alternative would place more equipment in one location, providing easier access for maintenance and diagnostic functions from a centralized location.

7.3. Synchronous Time-Division Multiplexing (STDM)

317



(a) Completely Centralized Demultiplex Approach



(b) Drop and Insert Approach, Where New York Subchannels Are Available in Chicago

Figure 7.7. Configuration options in a cascaded STDM network.

The user planning a cascaded STDM network needs to consider carefully whether the flow patterns in his network are centralized or noncentralized, and how the initial STDM network layout could be affected by a need to add new channels at a future date. The strategy of performing as much demultiplexing of inbound data streams as possible at the central site appears to offer numerous advantages in centralized-flow situations. On the other hand, the drop-and-insert scheme of Figure 7.7(b) affords more flexibility in decentralized applications where the subchannels need to be directly accessible at intermediate locations in the STDM cascade.

7.3.2. Similarities and Differences between FDM and STDM

Both FDM and STDM are widely used for reducing costs in end-user networks, the cost-reduction possibilities in either case arising from often-present economies of scale in the cost of bandwidth. Meaningful cost comparisons of various multiplexing techniques obviously require that necessary mileage-independent costs such as those for line terminations, multiplexers, and signal converters also be included. However, many present tariffs are structured so that multiplexing can produce substantial net savings, even after the costs of all required equipment are factored into the overall comparison. When the aggregate low speed bit rate for all terminals does not exceed 2000 bits/sec (give or take 10%), either FDM or STDM can probably be used; however, with current technology, FDM will probably be more cost effective than STDM, particularly when remote terminals are not clustered at a single site. Whenever a higher aggregate bit rate is required or any synchronous terminals are included in the sharing group, STDM will usually be dictated. However, in higher bit rate applications involving geographically dispersed terminal locations, an integrated blend of FDM and STDM will be the most sensible choice. Frequency-division multiplexing can span isolated terminal sites, creating traffic clusters that are then synchronously multiplexed into one or more computer sites.

Historically, the predominant usage of multiplexing has involved the derivation of low speed teletypewriter-grade channels on voice-grade lines. More recently, newer applications of STDM have appeared, particularly with the increased availability of higher speed synchronous modems, the initial availability of all digital data networks from conventional carriers, the entry of specialized carriers into the data network business, and recent provisions enabling customer-provided multiplexing equipment to be used over carrier-provided wideband links. If the costs of a multiplexed wideband link between two points can initially be justified, users may expect generally lower error rates on all derived voice-grade and low speed channels, substantially increased flexibility, and the opportunity to assign initially unused capacity at a later date without increased modem or line costs on the shared link.

These points are now illustrated in detail using specific examples. The reader is cautioned that the tariffs used in these examples are strictly for illustration purposes. Exact rates should always be obtained from the carrier. Tariffs used were in effect at the time when comparisons were made.

7.3. Synchronous Time-Division Multiplexing (STDM)

319

Example Problem 1: Comparison of FDM and TDM

Assume that a central computer located in downtown Chicago needs to provide 712 mile connections to 10 separate terminals in the same building in New York City. The terminals operate at 110 b/s/sec. Determine whether individual leased lines, frequency-division multiplexed analog lines, or time-division multiplexed analog lines would be the most cost-effective networking strategy. The cost assumptions below are illustrative of typical industry prices for comparable equipment at publication time.

Cost Assumptions:

Cost of FDM equipment: \$30 per month per channel end
 Cost of TDM equipment: \$250 per month (fixed) plus \$20 per month per low speed port
 Cost of modem equipment: \$50 per month for 2400 bits/sec units
 \$100 per month for 4800 bits/sec units
 \$200 per month for 9600 bits/sec units
 Cost of individual low speed lines—AT&T Series 1006 FDX
 Cost of voice-grade lines—AT&T Series 3002 (MPL)

Option 1: Individual Low Speed Lines

Monthly Cost of Individual 110 bit/sec circuit, including signal conversion equipment


$$= \underbrace{2.023 \times 100 + 1.416 \times 150 + .811 \times 250 + .605 \times 212}_{\text{Mileage}}$$

Mileage

$$+ \underbrace{2(36.15 + 14.45)}_{\text{Terminations}}$$

Terminations

So the monthly cost of 10 circuits provided with individual lines = \$8469.10.

Option 2: Frequency-Division Multiplexing Approach

320


Multiplexing and Concentration Techniques for Line Sharing

Since there are 10 FDM channel ends at each end of the voice-grade line, the FDM equipment cost is $20 \times 30 = \$600/\text{month}$. The monthly cost of the line (assuming no special conditioning) is

$$175.20 + 0.66(712 - 100) + 2 \times 25 = \$629.12$$

Hence the total cost of the FDM approach is \$1229.12. This is clearly more attractive than Option 1.

Option 3: Time-Division Multiplexing Approach



Monthly TDM costs are $2(250 + 10 \times 20)$	= \$900
Monthly modem costs are 2×50	= 100
Monthly line cost is (from above)	= <u>629.12</u>
TOTAL COST	\$1629.12

Hence the FDM approach is the most attractive of the three considered.

The reader should also consider other types of approaches such as WATS, dial-up, or the packet switching networks of value-added carriers before choosing a specific configuration. However, these options would require some idea of traffic volumes and usage patterns.

Example Problem 2: Another Comparison of FDM and TDM

Assume the same problem as for Example 1 except that the number of terminals in New York is increased from 10 to 20. Assume that an individual FDM system can derive a maximum of 12 subchannels at 110 bits/sec over one voice-grade line and that the TDM system can provide 20 subchannels on a 2400 bits/sec line, 40 subchannels on a 4800 bits/sec line, and 80 subchannels of 110 bits/sec on a 9600 bits/sec line.

Option 1:

$$\text{Monthly cost} = 20 \times 840.50 = \$16,810$$

Option 2:

$$\text{Monthly cost} = 2 \times 1229.12 = \$2458.24$$

(since two separate lines and four FDM devices are now required).

7.3. Synchronous Time-Division Multiplexing (STDM)

321

Option 3:

$$\text{Monthly TDM cost} = 2[250 + (20 \times 20)] = 1300.00$$

$$\text{Monthly modem cost} = 2 \times 50 = 100.00$$

$$\text{Monthly line cost} = 629.12$$


TOTAL COST

\$2029.12


Now the TDM approach is more favorable because of the required number of channels.

Example Problem 3: Synchronous TDM Possibilities

Assume that a Chicago computer center requires eight ports of 4800 bits/sec for connections of different synchronous terminals in a New York City regional office. Find the best way to provide the service, assuming that the following alternatives are available: (a) individual analog voice lines with modems, (b) wideband analog lines, and (c) digital service such as DDS (see Chapter 3 for service explanation and cost assumptions). Use the same equipment costs except that the cost of a synchronous TDM input port is assumed to be \$40/month.

Option 1: Individual Analog Lines

$$\text{Monthly cost} = 8 \times (629.12 + 200) = \$6632.96$$

Option 2: TDM over Wideband Analog Lines

Monthly line cost of AT&T Series 8000 wideband analog lines, including signal conversion equipment,

322

Multiplexing and Concentration Techniques for Line Sharing

$$= \underbrace{16.20 \times 250 + 11.40 \times 250 + 8.15 \times 212}_{\text{Mileage}} \times \underbrace{2 \times 460}_{\text{Terminations}} = \$9547.80$$

$$\text{Monthly TDM cost} = 2[250 + (8 \times 40)] = \$1140$$

$$\text{TOTAL COST} = \$10,687.80$$

*Option 3: Individual Dataphone Digital Service (DDS)**Lines at 4800 bits/sec*

Monthly cost of individual 4800 bit/sec DDS line

$$= \underbrace{0.26 \times 71}_{\text{Mileage}} + \underbrace{2 \times 20.60}_{\text{Intercity link termination}} + \underbrace{2 \times 87.55}_{\text{Local access line}} + \underbrace{2 \times 15.45}_{\text{Signal conversion}} = 688.64$$

$$\text{Monthly cost of eight separate lines} = \$5509.12$$

Option 4: Multiplexed 56,000 bits/sec Dataphone Digital Service

Monthly line cost, including signal conversion equipment,

$$= 4.12 \times 712 + 2 \times 64.50 + 2 \times 206 + 2 \times 20.60 = \$3515.64$$

$$\text{Monthly TDM cost (from above)} = \$1140$$

$$\text{TOTAL COST} = \$4655.64$$

Summary of Example

The best solution for this example appears to be multiplexed 56,000 bits/sec Dataphone Digital Service. However, the monthly cost savings over individual 4800 bits/sec DDS lines needs to be weighed against the fact that the reliability properties of the multiplexed network are much poorer. A failure of the TDM equipment or of the multiplexed line would cause all channels to become unavailable. The operational costs of such a catastrophic failure situation may mean that the multiplexed network is less desirable, in spite of the cost savings it offers.

The reader is also cautioned against generalizing about the relative merits of various networking approaches from this example. The conclusions here, and for other examples as well, are strongly dependent on all the cost assumptions, distances, and required numbers of channels. However, the solution techniques are quite general and are equally useful for alternative tariff structures and equipment costs.

7.4. STATISTICAL TIME-DIVISION MULTIPLEXING (STATDM)

Much of the theoretical work relating to STATDM has been done by Chu [2,5-8]. Related contributions have been made by Rudin [20], Birdsall et

7.4. Statistical Time-Division Multiplexing (STATDM)

323

al. [9], Pan [10], Gordon et al. [11], and Chang [12]. Reference [30] also discusses STATDM in detail. Statistical time-division multiplexing differs from STDM in that a dedicated subchannel is not provided for each port in the sharing group. Since, under certain conditions of heavy loading, a STATDM may be incapable of accommodating all the terminals in its sharing group, statistics and queuing become important considerations. Thus it is a hybrid form of multiplexing and concentration.

The fundamental idea of STATDM is to exploit the property of STDM systems that many of the time slots in the fixed-format frames are wasted because a typical sending terminal will actually be transmitting data less than 10% of the time it is communicating with the CPU. A more detailed discussion of typical traffic arrival statistics is presented by Fuchs and Jackson [13]. As shown in Figure 7.8, STATDM dynamically allocates the time slots in a frame of data to the currently active users, reducing the fraction of wasted time slots and thereby increasing overall line utilization and throughput.

Although the diagram of Figure 7.8 illustrates addressing information being transmitted with data in each slot, it is of course not necessary to do so in cases where such a procedure could lead to excessive overhead. An alternative would be to send demultiplexing address information in a control frame only once at the beginning of each dynamic subchannel establishment. This demultiplexing rule can be dynamically updated only when subchannels are added or removed, without the need to include address information bits explicitly with the data in each slot. Another possibility is to vary the slot widths for the individual ports or to encode a control signal that tells the demultiplexer exactly which ports are idle in a given frame.

Most estimates of the exact performance improvements attainable with STATDM over STDM have to date been based on analytical studies described in certain of the references previously cited. From Chu's analytical studies, it would appear that from two to four times as many users could be accommodated on a voice-grade line as with STDM, assuming an application environment where either method could be used. In certain situations where low duty cycle terminals are serviced by a statistical multiplexer over a broadband link, the margin could be substantially greater.

The tradeoff disadvantages of statistical multiplexing are the costs of substantially more elaborate addressing and control circuitry, the need for data buffers to hold incoming messages, and the possibility of blocking and/or appreciable queuing delays under heavily loaded conditions. The references previously noted contain substantial traffic studies, investigating the relationships between such factors as traffic intensity, distribu-

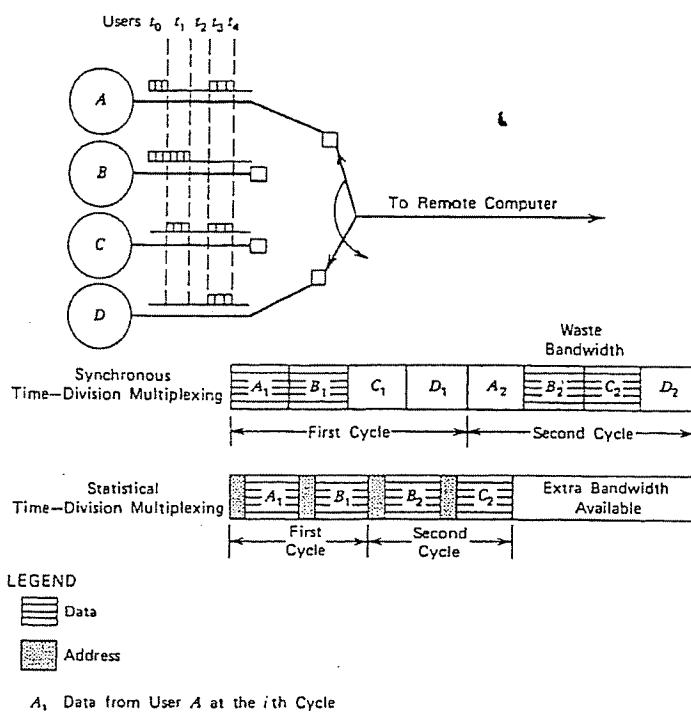
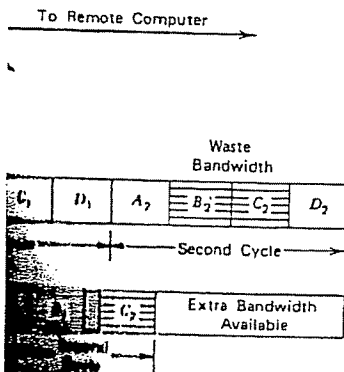


Figure 7.8. STDM contrasted with STATDM.

tions of message arrivals and lengths, queue sizes, queuing discipline, and blocking probability. To illustrate, several of the major results described in Chu [2,5-8] are now summarized.

Chu has analyzed a Markov model of a statistical multiplexer in which messages arrive at a finite-capacity multiplexer according to a batched Poisson process where the size of the batch corresponds to the length of the arriving message in characters. A unit service interval μ is the time to transmit a character on the shared line; for a synchronous line with a transmission speed of R characters/sec. $\mu = 1/R$. Message lengths are assumed to be geometrically distributed with mean \bar{l} , and the number of messages arriving during a unit service interval is Poisson distributed with a mean rate of λ messages every 1 sec. The buffer overflow probability is obtained from the steady-state solution to the state equations for an embedded Markov chain. The average queuing delay per message D is shown to be given by

Concentration Techniques for Line Sharing



7.4. Statistical Time-Division Multiplexing (STATDM)

325

$$D = \frac{\lambda(2 - \theta)}{2(\theta - \lambda)\theta} \quad (\text{character service times})$$

where $\theta = 1/\bar{l}$. The buffer overflow probability is assumed to be sufficiently small that virtually all traffic arriving at the multiplexer is transmitted over the line.

At the demultiplexing end of the line, Chu has used a simulation model to estimate overflow probabilities and the following analytic relationship to describe average waiting time W_i for sending messages to the i th destination:

$$W_i = \frac{\rho_i(2\bar{l}_i - 1)}{2(1 - \rho_i)} \quad (\text{character service times})$$

where \bar{l}_i = average message length for the i th destination

λ_i = message arrival rate for the i th destination

μ_i = transmission rate for the i th destination

$\rho_i = \lambda_i \bar{l}_i / \mu_i$

Figure 7.9 indicates the parameters of a generalized model that may be used to evaluate design tradeoffs in configuring statistical multiplexers. It also suggests that a statistical multiplexer can perform two levels of concentration when the input terminals are not permanently connected to the input ports.

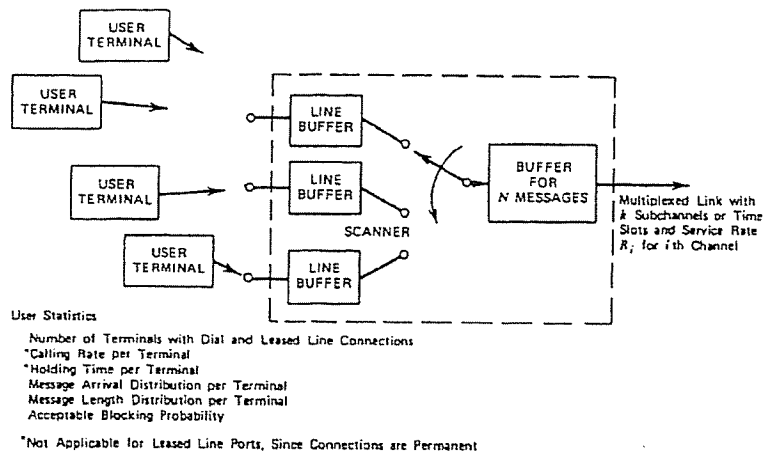


Figure 7.9. Schematic diagram of STATDM for traffic studies.

It appears that STATDM has a promising future, particularly in applications where queuing delays are not of material concern or can be readily minimized. It is this author's conviction that statistical multiplexers will be of primary use, not in replacing STDM en masse, but in new

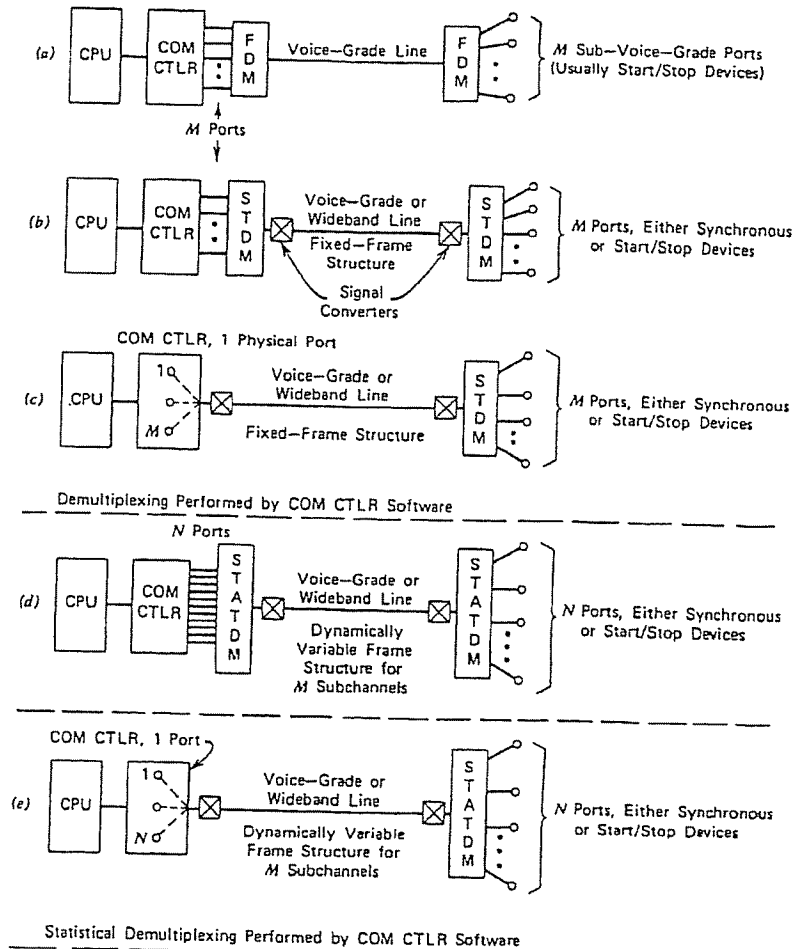


Figure 7.10. FDM, STDM, and STATDM approaches to line sharing. In the STATDM approaches, the number of physical ports, N , connected to the sharing devices is usually larger than M , the number of subchannels created at any instant in time. M changes with time because of fluctuations in user traffic patterns.

7.5. Message and Packet Switching Concentration

327

applications involving store-and-forward message switching, loop transmission systems, system-provided error control, and so on. For example, CRT display controllers can be equipped to implement STATDM in conjunction with ARQ error control. A request for retransmission would be issued either when a data block error is detected at the receiving terminal or when the buffer area in the multiplexing control unit is full. One of the major problems is incorporating enough intelligence in STATDM to anticipate when an input line is about to become active so that proper steps may be taken to assign the next available time slots to the user in question. A related problem lies in accurately sensing when a terminal has completed its transmission so that time slots are not filled with "idle" characters. In certain applications this problem can be mitigated somewhat by having the STATDM unit track the input buffers for special end-of-message or end-of-block characters. However, this approach requires a knowledge of the terminal code format and would be of limited use in transparent text applications.

Figure 7.10 illustrates the relative equipment requirements for using FDM, STDM, and STATDM to service a remote cluster of terminals. Note that demultiplexing of inbound channels may take place either in a stand-alone box or in the communications control unit. Because the STATDM approach requires user traffic flows to be monitored anyway, the combined-function, single-box approach makes the most conceptual and economic sense. The only potential problem with the combined-function approach is the issue of communications control unit horsepower. Since STATDM requires virtually continuous tracking of the ports, an already heavily loaded communications control unit may not be able to accommodate the added STATDM function. Even with contemporary microprocessor technology, a separate STATDM unit may often be necessary at the central site.

Another substantial advantage of STATDM over STDM is its flexibility in providing different subchannel mixtures across the shared link at different times. This benefit is essentially independent of the increases in throughput resulting from statistical use of shared lines. For example, an STATDM system could easily function as an STATDM for a while and then convert automatically to operate as a conventional STDM system with different subchannel mixtures at other times.

7.5. MESSAGE AND PACKET SWITCHING CONCENTRATION

Message switching concentration (MSC) and packet switching concentration are functional extensions of statistical multiplexing involving the "multiplexing" of entire messages and fixed-length portions of long messages, respectively. They are more properly categorized as concentration

EXHIBIT 18

82-0454

MULTIPLEXING OF PACKET SPEECH ON AN EXPERIMENTAL WIDEBAND SATELLITE NETWORK

C. J. Weinstein and H. M. Heggstad

Massachusetts Institute of Technology
Lincoln Laboratory
Lexington, Massachusetts 02173

ABSTRACT

Packet techniques provide powerful mechanisms for the sharing of communication resources among users with time-varying demands. In particular, significant potential advantages are expected for systems which integrate digital data and digitized voice in a common packet-switched network, including bandwidth savings and improvements in channel utilization and flexibility. An experimental system under development by DARPA, consisting of a wideband demand-assigned satellite network and a variety of local access networks with associated gateways, is being used to explore these expected characteristics. The components of the system are described in detail, and progress and goals in the experimental program are described.

I. INTRODUCTION

Packet techniques provide powerful mechanisms for the sharing of communication resources among users with time-varying demands. Because of their discrete and independent nature, outgoing packets can be inserted in any available time interval on a communication channel. By sharing the channel among a pool of users who each maintain a queue of packets ready for transmission, one can trade a moderate increase in overall delay for a high utilization of the channel. Packet satellite systems [1] in particular, provide the opportunity for efficient demand-assigned sharing [2,3] of long-haul communications capacity among users at widely dispersed locations. The primary application of packet techniques has been for digital data communications, where the bursty nature of user traffic can be exploited to achieve large efficiency advantages in utilization of communication resources.

Packet techniques offer significant potential advantages for voice as well. The integration of digital voice and data into a common network offers potential cost advantages [4] associated with resource sharing among different types of users, as well as the possibility for enhanced services for users who require access to both voice and data communications. Packet networks offer advantages for digital voice conferencing [5] in terms of channel utilization and control flexibility. Channel capacity savings for packet voice are achieved by transmitting packets only when speakers are actually

talking (i.e., during "talkspurts"). Then, by multiplexing many packet speech users on a wideband channel a factor approaching two in bandwidth saving can be achieved, similar to that achieved on specific long-haul circuit-switched network links through Time-Assigned Speech Interpolation [6] (TASI) or Digital Speech Interpolation [7] (DSI). Other benefits are obtained with packet voice. A packet network allows convenient accommodation of voice terminals with different data rates and data formats. If vocoders are used to reduce bandwidth utilization, then each vocoder will use only the capacity necessary to transmit its information rather than the fixed minimum bandwidth increment typically used in circuit-switched networks. Packet networks also provide a convenient system environment for implementation of voice flow control techniques [8] where the bit rate is dynamically adapted to network conditions.

An experimental wideband satellite-based packet system [9] is being implemented to develop and demonstrate techniques for achieving the advantages of integrating packet voice with data in a realistic large scale network. The research supported by this system builds upon previous experiments on the ARPANET [10,11] and on the Atlantic Packet Satellite Network (SATNET) [12]. In those experiments the basic feasibility of packet speech was demonstrated, techniques were developed [10,13] for reconstitution of speech from packets arriving at non-uniform intervals, and packet protocols [14] were developed for call set-up and for speech transport. However, the network link capacities were too small to accommodate a sufficient number of simultaneous users to allow experimental demonstration of efficient statistical multiplexing of voice. The new experiment is designed around a satellite channel with a capacity of 3 Mbps, which will be able to support many simultaneous voice connections. In addition to a practical demonstration of efficient statistical multiplexing, a key purpose of the experiment is to assess the implementation requirements for terminals and for switching and multiplexing in a large-scale packet speech system.

The experimental system consists of research facilities at multiple sites which are linked by a wideband packet satellite network (the NR SATNET) including a channel on a satellite transponder, earth stations, high-performance burst modems, and demand

assignment multiple access (DAMA) processors. The sites have local packet speech access facilities including packet voice terminals [15] (PVTs), local area networks including cable [16,17] and radio [18] systems, and concentrators [19,20]. The overall system design has a high degree of flexibility, so that a variety of experiments can be supported. Besides access facilities for real traffic, the system includes traffic emulators to allow controlled variation of network loading. Measurement facilities are provided, to support experimental evaluation of system performance.

The wideband packet speech system development and the experiment program are sponsored by Defense Advanced Research Projects Agency (DARPA), and involve a cooperative effort by a number of organizations. These include: Bolt Beranek and Newman, Inc. (BBN), Cambridge, MA; COMSAT Laboratories, Clarksburg, MD; Information Sciences Institute (ISI), Marina del Rey, CA; LINKABIT Corp., San Diego, CA; MIT Lincoln Laboratory (LL), Lexington, MA; SRI International, Palo Alto, CA; and Western Union, Inc., Upper Saddle River, NJ. The Defense Communications Agency (DCA) has supported the development of the satellite network along with DARPA, and is utilizing the WB SATNET for a set of experiments supporting the development of the future Defense Communications System. One of the four presently-installed network nodes is located at the Defense Communications Engineering Center (DCEC) in Reston, VA.

This paper describes the experimental wideband packet system and the ongoing experiments in packet speech multiplexing. Section II presents a system overview, and describes how the WB SATNET fits into a context of an interconnection (or internetwork) of local and long-haul packet nets. Specific subsystems of the WB SATNET, and design considerations underlying their development, will be described in Section III. Priority-Oriented Demand Assignment (PODA) [1], the demand assignment multiple access (DAMA) algorithm used for efficient sharing of the satellite channel, is reviewed in Section IV. Section V describes the subsystems implemented for multiplexing of local area traffic, and Section VI discusses the packet voice protocols used for concentration of speech traffic. Section VII reports on current and planned experimental activities, and a summary and conclusions are given in Section VIII.

II. SYSTEM OVERVIEW

Currently there are four WB SATNET sites, located at DCEC, ISI, LL, and SRI. At each site the WB SATNET equipment includes an earth station, burst modem, and packet satellite DAMA processor. These elements are included in the region labeled WB SATNET in Fig. 1, and will be described in Section III. The burst modem and DAMA processor are designed with a high degree of flexibility so that a variety of experiments can be supported. The WB SATNET performs the task of accepting packets from external sources, managing the satellite channel allocation on a demand assigned basis, and delivering the packets to their

destination. As indicated in Fig. 1, the WB SATNET provides long-haul broadcast connectivity among several local area networks. At LL and at ISI there are local broadcast cable networks (referred to as LEXNET for Lincoln Experimental Network) which were developed at LL to efficiently support local packet voice and data traffic. LEXNET uses a carrier-sense multiple access protocol with collision detection (CSMA-CD) similar to that used in ETHERNET [21], but includes a specialized distributed access control mechanism [16], designed to take advantage of the fact that speech traffic contains bursts of periodically-generated packets. At SRI the "local" net is a packet radio network [22] which includes provision for mobile users as well as users at fixed positions.

Access from the local area onto the WB SATNET is obtained by means of an internetwork gateway at each site, as shown in Fig. 1. The gateways carry out a number of tasks including (1) communication with the WB SATNET nodes; (2) inserting the specific WB SATNET headers needed to transport long-distance packets to the destination gateway; (3) concentrating speech packets from a number of local terminals into aggregated packets for the satellite net; and (4) requesting satellite capacity allocation on the WB SATNET based on rate requirements identified by packet voice terminals at dial-up. The channel allocation requirement is ideally set on a statistical basis and takes account of the fact that voice packets are transmitted only during talkspurts. This resource allocation function for gateways is part of the "stream" or ST protocol [23,19] which is implemented experimentally in the wideband system, and is compatible with the DoD standard internet data packet (IP) protocol (see [24]). IP gateways treat packets as independent datagrams; ST gateways provide real-time stream allocations as well as packet aggregation.

The internetwork configuration shown in Fig. 1 also includes gateway connections to the ARPANET, a long-haul store-and-forward packet-switching network connecting research facilities across the U.S. and in Europe. This connection allows intercommunication among computers (referred to as "hosts," on the ARPANET and terminals and hosts in the other local networks. For example, the network control center (NCC) for the WB SATNET resides in an ARPANET host computer at BBN. Direct ARPANET access is provided from this BBN host computer to the DAMA processors in the WB SATNET, supporting remote monitoring and control functions which are necessary during installation and checkout as well as for some categories of experiments. Voice communication can be carried out between experimental packet voice facilities developed for the ARPANET and new facilities on LEXNETs and PRNETs. On a limited basis, terrestrial routing via the ARPANET (e.g., from LEXNET 1 to LEXNET 2) can be an alternative to satellite routing, provided that sufficient voice bit rate compression is applied so that

the limited throughput capability of the ARPANET can handle the real-time transmission rate.

Two kinds of interfaces are shown in Fig. 1 between the packet-switched system and circuit-switched systems of the type more commonly used for voice communication. Special circuit/packet interfaces have been developed [25] for the gateways at LL and DCEC to allow communication with a digital circuit switch in the T1 digital carrier format used for interswitch communication in digital telephony. This allows communication between voice terminals on circuit-switched and packet-switched networks, but more importantly provides the capability for experiments in which a DAMA satellite network is used as an overlay to a terrestrial circuit-switched net. These satellite overlay experiments are being implemented, under DCA sponsorship, to develop networking techniques applicable to the next generation CONUS AUTOVON [26] and a future Defense Switched Network (DSN) [27] which will utilize a mix of transmission media to provide survivable and economic digital voice and data service to DoD subscribers.

The switched telephone network (STN) interface developed at ISI allows connection between individual telephone lines (rather than switches) and packet voice terminals (PVTs) on a dial-up basis. This interface allows a user to dial into the packet voice system from any remote telephone. Plans also call for STN interfaces to be installed at other sites in the system.

With this system overview as background, the next section will proceed with a more detailed description of the functions and subsystems of the wideband packet satellite network.

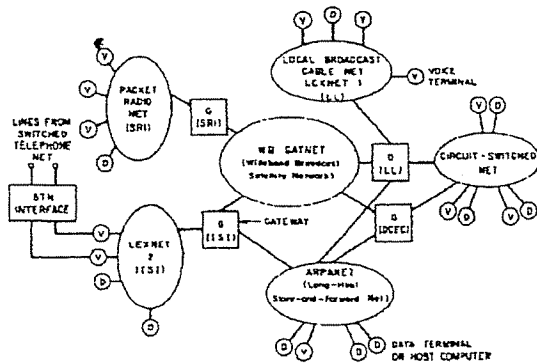


Fig. 1. Integrated voice/data packet internetwork.

III. PACKET SATELLITE NETWORK SUBSYSTEMS

The WB SATNET provides packet-switched broadcast satellite connectivity among the local networks at the four experimental sites (DCEC, ISI, LL and SRI). The major WB SATNET subsystems required at each site, illustrated in Fig. 2, are the antenna and earth station; the Earth Station Interface (ESI); and the Pluribus Satellite Interface Message Processor (PSAT). The Lincoln Laboratory site configuration, including the WB SATNET subsystems as well as a concentrator and a local network, is typical of the wideband experiment sites; it is illustrated in Fig. 3.

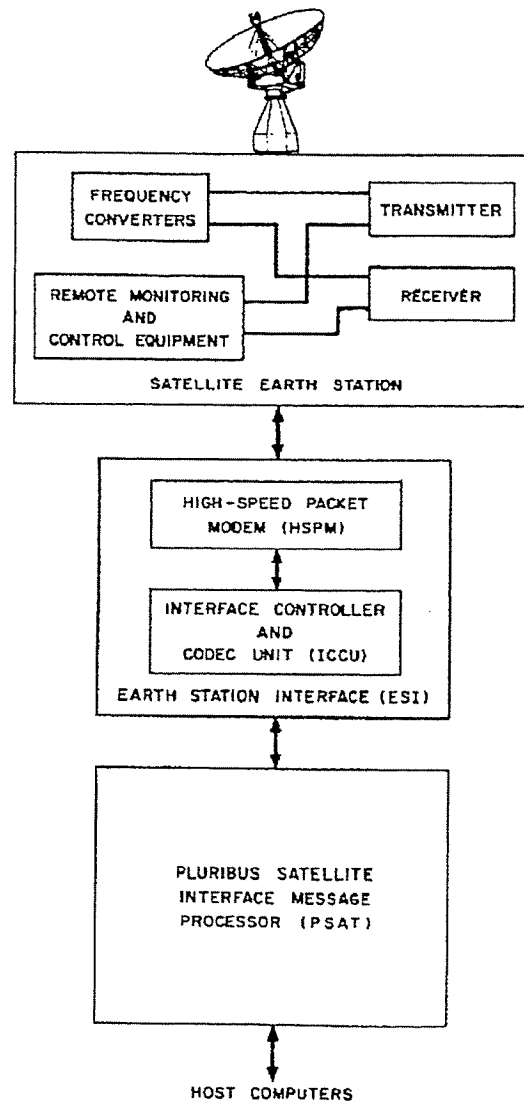


Fig. 2. Block diagram of wideband packet satellite network subsystems.




Fig. 3. Typical site configuration for experimental wideband packet network.

The satellite link, antenna and earth station subsystem are leased from Western Union, Inc. Satellite capacity is provided in the 4/6 GHz commercial telecommunications band. The WB SATNET carrier is at the upper end of Transponder 1 on Western Union's WESTAR III satellite, which is located at about 91 degrees west longitude. The exact frequencies are 5.959 GHz (uplink) and 3.734 GHz (downlink), and the available bandwidth is matched to the 3.088 Mbps digital signal used by the system. The antenna (a 5-meter Cassegrainian) has a field-effect transistor (FET) receiver front-end amplifier mounted behind its feed horn; together they provide an earth station figure of merit G/T of about 22 dB/Kelvin. Uplink and downlink power budgets are designed to support system operation at a specified maximum bit error rate of 5×10^{-6} at a channel bit rate of 3.088 Mbps. The transmitter amplifier is a 125-watt traveling-wave tube (TWT); it is installed in a small shelter at ground level behind the antenna in Fig. 3, along with suitable frequency up- and down-converters and IF cable drivers and receivers. In order for the equipment to be operated as an unmanned earth station, it is provided with remote monitoring and control facilities connected via automatic telephone dialing equipment with a central satellite control facility in New Jersey.

The earth station interfaces at 84 MHz via analog IF cables with a flexible burst modem (the High-Speed Packet Modem or HSPM). This subsystem together with the Interface Controller and Codec Unit (ICCU) constitutes the Earth Station Interface (ESI) equipment, which is manufactured by LINKABIT. The HSPM offers a choice of binary or quaternary phase-shift keying (BPSK or QPSK), at channel symbol rates of 193, 386, 772, or 1,544 kilobaud. The bit rate with QPSK is twice

the channel symbol rate, and hence the maximum bit rate over the channel is 3.088 Mbps. The Codec implements convolutional encoding and sequential decoding at a selection of code rates (1, 7/8, 3/4 and 1/2) with corresponding coding gains equivalent to signal-to-noise ratio improvements of 0 dB to as much as 5 dB. Suitable combinations of bit rate and code rate can be chosen to accommodate a broad range of bit error rate tolerance levels; for example, TDMA framing and synchronization information and packet headers must be received very accurately, and hence should be more heavily protected with coding (i.e., a low code rate should be used). The portions of packets which contain speech signals can tolerate increased bit error rates, on the other hand, and therefore permit operation at a higher code rate with a corresponding gain in channel throughput. Reduced code and symbol rates can also be chosen to permit interoperability with less sensitive earth stations, having smaller antennas or higher receiver noise temperatures. An important flexibility feature of the ESI is that it supports multiple changes in code rate and bit rate within a TDMA burst. Implementation of burst formats and control of the functions of the ESI are accomplished in a Motorola MC68000 16-bit microprocessor in the ICCU.

The format and contents of the transmissions in each TDMA frame are supplied to the ESI by the Pluribus Satellite Interface Message Processor, or PSAT [28]. The PSAT is a high-throughput multi-processor computer built by Bolt Beranek and Newman. The primary functions of the PSAT are management of demand-assigned allocation of time slots on the satellite channel, and transmission and reception of information bursts over the channel. The Pluribus architecture combines the processing power of six Lockheed SUE processors. Satellite demand assignment is arbitrated by a distributed protocol which combines features of both circuit- and packet-switched control techniques. This protocol, which is called the Priority-Oriented Demand Assignment (PODA) algorithm [1,28,29], is discussed in the following section.

IV. SATELLITE CHANNEL DEMAND ASSIGNMENT ALGORITHM

The Priority-Oriented Demand Assignment (PUDA) algorithm was developed for the Atlantic Packet Satellite Experiment [30] (the Atlantic SATNET) which, as noted in Section I, links a number of earth stations in the United States and Europe by means of a 64 kbps INTELSAT satellite channel. The implementation of PUDA for the WB SATNET, which includes a number of modifications to adapt to the higher bit rate of the channel, is described in detail in [28].

The PSAT's transmissions basically constitute a Time-Division Multiple Access (TDMA) structure with variable frame format and fixed frame duration. A dynamic channel-sharing discipline such as PODA undertakes to achieve high efficiency on the channel for

the relatively low-duty-cycle traffic which is characteristic of WB SATNET users by dynamically re-allocating unused capacity among the subset of users who have traffic bursts to send at any given time. Traditional fixed-format TDMA and Frequency-Division Multiple Access (FDMA) are static channel-sharing disciplines which permanently allocate capacity among the community of users. Ideally, these capacity allocations are sized to match the users' communication requirements; in practice the efficiency of utilization of the channel can be quite low, however, because random variations in user activity tend to leave the assigned blocks of capacity empty some fraction of the time. A user's data stream can be communicated very efficiently using fixed-format TDMA or FDMA if it has a high duty cycle, corresponding to (for example) multiplexed signals from a large local population of voice or data terminals. A user data stream containing multiplexed traffic from a moderate to small local population of terminals will result in bursty traffic with a lower duty cycle, which is best handled by a dynamic technique such as PODA.

Key functions that must be accommodated by a dynamic channel-sharing discipline include user capacity request mechanisms, determination of assignments in response to the requests, and achievement of low assignment delay coupled with efficient channel usage. Additional functions that would further enhance a dynamic discipline include the ability to adapt to both error-tolerant traffic (e.g., speech signals) and error-intolerant traffic (e.g., computer file transfers), and the ability to accommodate both strong stations and weaker stations (e.g., those which have smaller antennas or less sensitive receivers, or are temporarily suffering reduced performance due to rainfall attenuation or component degradation). The manner in which PODA incorporates these functions is described briefly below.

PODA provides two mechanisms for user capacity requests: use of a special reservation subframe, and "piggybacking" new requests upon ongoing transmissions. Two types of reservation subframe techniques have been implemented: fixed PODA (FPODA), in which each user is permanently allocated a specific request slot, and contention PODA (CPODA), in which request slots are shared among users on a contention basis. PODA is presently implemented as a distributed-control discipline. Each PSAT hears all capacity requests, and all the PSATs arrive at the same set of slot assignments by applying the identical algorithm to the set of requests. This results in a two-hop delay for transmissions from any given user (one satellite round trip time for submitting the capacity request, and one round trip time for actual transmission of the information in the assigned slot). Centralized assignment, in which one processor computes slot assignment for all users, results in a three-hop delay because of the extra round trip time required for the central controller to communicate its slot assignment decisions back to the requesting stations. The PODA design includes centralized assignment for a limited subset of users,

namely weak stations which cannot hear all of the transmissions exchanged across the network. This centralized assignment capability is planned but not yet implemented. The PSAT can accommodate weak stations by a combination of the centralized assignment feature of PODA and the capability to make suitable code and bit rate selections by commands to the ESI.

Figure 4 is a generalized illustration of a PODA frame (typically 21.2 msec in duration in the present implementation). The information and control subframes are used as described above. Within the information subframe there are sequences of information bursts, each from a PSAT which has obtained a slot assignment in that frame. An information burst can consist of either datagram packets or concentrated stream packets. Datagrams are independent self-contained messages which are communicated between concentrators and the PSAT in accordance with the DoD standard internet data protocol (IP); they are used for delay-tolerant digital data communications such as computer-to-computer file transfers, and can offer error-free service through the use of mechanisms such as acknowledgement and retransmission. Stream packets are elements of extended real-time communication sessions such as digitized speech streams, which have the properties that (1) they will tolerate little delay, and little or no delay variation; (2) once initiated, they may be expected to continue for a time interval extending over many PODA frames. Streams are handled in accordance with a new experimental protocol called ST, which is discussed further in Section VI. PODA accommodates streams by offering stream reservations (initiated by means of requests made by the concentrator to the PSAT at call set-up time) for an extended sequence of packets providing transmission for the stated average bit rate for the duration of the call. The bit error tolerance of voice avoids the need for re-transmission, thereby tending to stabilize the delay. The gateways shown in Fig. 1 enhance the efficiency of satellite channel utilization by providing the PSAT with concentrated stream packets which have been subjected to destination-oriented aggregation, as explained in Section VI.

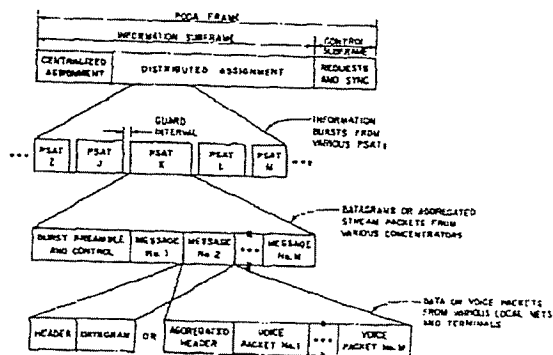


Fig. 4. Frame structure for Priority-Oriented Demand Assignment (PODA).

V. LOCAL AREA NETWORK SUBSYSTEMS

A block diagram of the local area network subsystems for a typical site is shown in Fig. 5. The concentrator/gateway has been developed by LL with initial application as a LEXNET/NR SATNET traffic concentrator. It has been augmented with the required network interface software and hardware to serve (in different configurations) as a gateway to the ARPANET, PRNET, and a digital circuit-switched net as well. The hardware consists of a central gateway processor (Digital Equipment Corporation PDP-11/44 minicomputer), and network interface processors (UMC-Z80 universal interface computers produced by Associated Computer Consultants) and special hardware interfaces for each attached network. Since the bulk of the gateway traffic is to be forwarded, the data portions of packets need not be examined by the gateways. The gateway design also provides for an additional memory to be accessed by the Z80 processors on a multiport basis. This shared memory is intended to store the data portions of packets so that the load on the PDP-11 UNIBUS and memory will be minimized.

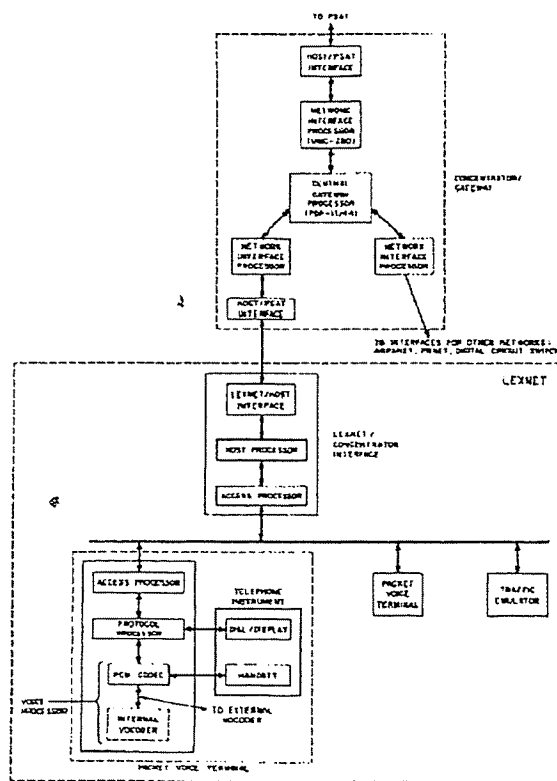


Fig. 5. Block diagram of local area network subsystems.

A functional block diagram for the concentrator is shown in Fig. 6. ST gateway functions deal with the processing of stream (i.e., voice) traffic, including the aggregation

of local area voice packets for efficient multiplexing on the WB SATNET. IP gateway functions handle datagram traffic, including the control packets needed to set up voice calls between packet voice terminals. Other functions include conference access control and control and monitoring. As indicated, the PDP-11 is the central gateway processor and the UMC-Z80s perform front-end packet handling tasks for each network. The NR SATNET UMC-Z80 performs the additional task of deaggregating ST packets destined for different terminals on LEXNET.

The LEXNET [17] is a baseband CSMA cable network with distributed control similar to ETHERNET [21]. It utilizes a distributed algorithm for randomized retransmission specialized for voice traffic [16]. This algorithm estimates competing network activity based on the fact that a voice terminal produces periodic packets during talkspurts, in contrast to data terminals where the inter-packet intervals are usually modeled as being independently distributed. The range of the retransmission interval after collisions is adjusted along with this activity estimate, with longer intervals corresponding to higher activity estimates.

The Packet Voice Terminal (PVT) [15] has been developed with a flexible and modular architecture, as illustrated in Fig. 5. There are three primary functional units, each controlled by a microprocessor. The voice processor digitizes the speech but is independent of the transmission medium. Each PVT has a 64 Kbps PCM capability and a selectable option of internally mounted vocoders of various rates, or connection to an external vocoder. The protocol processor forms the voice packets and provides the necessary layers of Network Voice Protocol (NVP) [14,19] to insure that the packet can be delivered to a distant network and played out at the proper time. The necessary buffering and speech reconstitution algorithms to produce steady speech to the listener despite asynchronous packet arrivals from the network are implemented in the protocol processor. The protocol processor also includes an interface with the user dial/display and must generate and interpret the packets necessary for establishing the call. A generalized packet interface is provided between the protocol processor and the access controller, which performs LEXNET-specific local net packet transport functions. It is possible to connect the PVT to other packet networks such as the PRNET by changing the design of the access module.

In addition to the PVTs, a traffic emulator is also provided in the experimental network to allow performance measurements to be made on the network with controlled amounts of traffic. The LEXNET traffic emulator module actually consists of a PVT with special protocol processor software to generate controlled traffic.

Packets entering or leaving LEXNET pass through the LEXNET/Concentrator Interface (LCI), a unit which is very similar to the

PVT. The access processor in the LCI is identical to that in the PVT, and the LCI host processor controls packet forwarding to and from the concentrator.

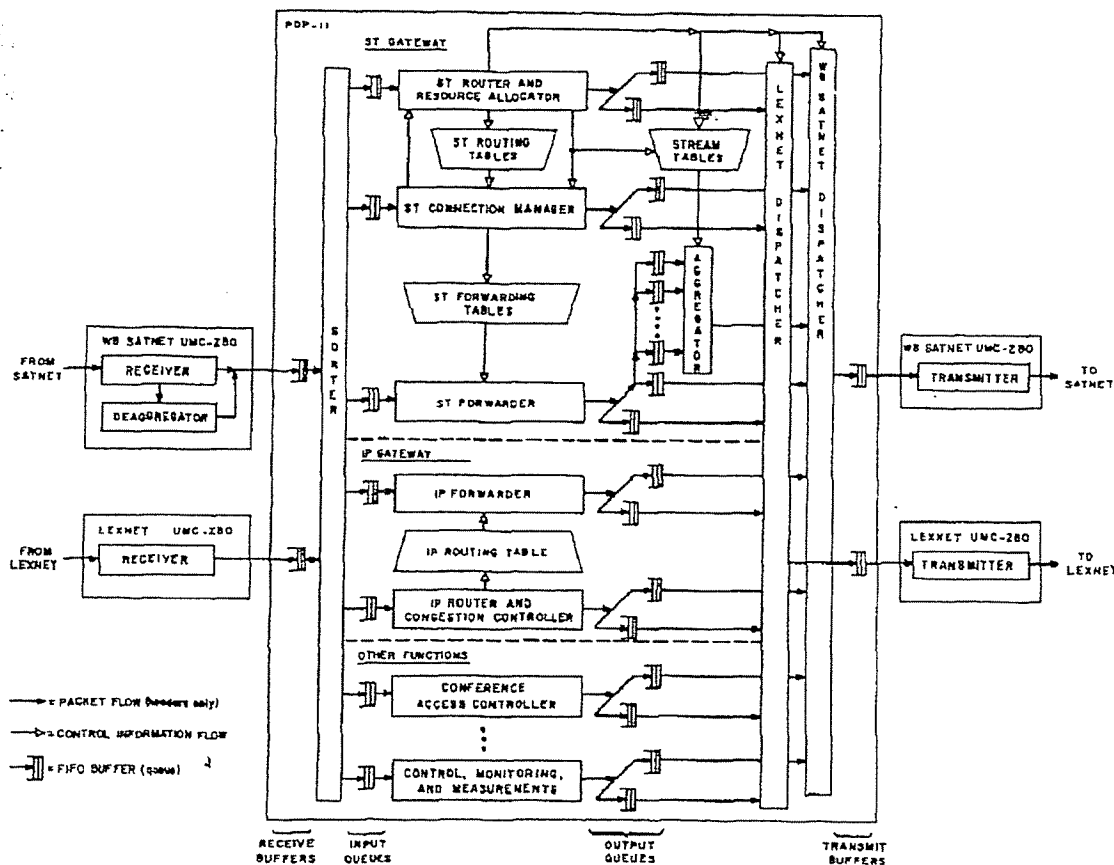


Fig. 6. Functional diagram of speech concentrator.

VI. PACKET SPEECH CONCENTRATION PROTOCOL

The format for an aggregated stream (ST) packet, as would be delivered by a concentrator to a PSAT to be transmitted in a stream slot allocated within the PODA frame, is illustrated in Fig. 7. The sample packet includes voice for two point-to-point (PTP) calls and one conference (CONF) call. All headers are grouped at the start of the packet. This allows advantage to be taken of the flexible coding properties of the burst modem (see Section III), by coding the headers (where bit errors are extremely damaging) more heavily than the speech (where some bit errors can be tolerated). The individual PTP headers are quite short, since they need contain only a connection identification

(CID) and not the full destination address. The ST connection process includes the set up of tables in gateways which associate CIDs with the next forwarding address along the route. The CONF header includes additional information to control the forwarding of packets to all conference participants. The data portion of the packets include a short (32-bit) MVP header with a time stamp to allow the PVT to determine the proper packet playout time. The number of speech bits depends on the voice bit rate and the interpacket interval. Typical ranges are 2.4-64 Kbps for voice bit rate, and 20-50 msec for the interpacket interval.




Fig. 7. Format for aggregated stream (ST) packet.

VII. CURRENT AND PLANNED EXPERIMENTAL ACTIVITY

A preliminary experiment plan [31] was prepared soon after the inception of the experimental wideband packet system program. Lincoln Laboratory has since been responsible for experiment planning and coordination for the overall program. Central to the experiment plans has been the notion of multiplexing sufficiently large numbers of packet speech users onto a common satellite channel to permit testing of the multiplexing efficiency. The experiments follow a logical sequence which has been keyed to achievement of successive stages of experimental capability, as the implementation of the various subsystems is completed.

The major experimental activity which has been completed at this writing has involved integration and performance demonstration in the two primary subdivisions of the system, namely the WB SATNET and the various local access facilities. The system is at the point of beginning to support the overall goal of the program, which is to carry out a broad range of multi-user packet speech experiments. These will include interconnections among voice terminals within local networks; among voice terminals on similar but separate

networks, internettted with each other via the WB SATNET; and among different but compatible types of voice terminals on dissimilar local networks internettted via the satellite net.

In the area of local access network integration and performance demonstration, the checkout of Packet Voice Terminals (PVTs) on LEXNETs has been substantially completed. Demonstration of two simultaneous conversations involving four PVTs on the same LEXNET was accomplished some time ago. Conversations between two PVTs on the same LEXNET have been looped back through the PSAT, FSI, and satellite transponder. A number of experiments have also been done in which PVTs on two separate LEXNETs have communicated with each other by way of a concentrator. Compatibility of the voice terminal and concentrator protocols at different sites (i.e., LL and ISI) has been established via narrowband (2400 bps) packet speech experiments using the ARPANET to link the distant local access facilities.

In the area of WB SATNET integration and checkout, the main thrust has been to achieve reliable operation of a network of PSATs performing their timing, synchronization and demand assignment functions via the satellite channel. This activity has been substantially completed with respect to a two-PSAT network (namely, those at ISI and LL).

The immediate goal at this time is to demonstrate multiple simultaneous cross-country conversations. Competing emulated traffic will then be introduced, and its intensity will be increased until system performance breakpoints can be observed. A broad range of variations on these experiments will be conducted, thus making it possible to evaluate the utility of packet switching for handling real-time multiple-user speech communication.

VIII. SUMMARY AND CONCLUSIONS

The background and motivation for experimentation with multi-user packet speech were described; cost advantages are anticipated, as well as improvements in flexibility and transmission channel utilization efficiency. The experimental wideband satellite-based network which is being implemented to support such experiments was described in some detail, including overall system design issues as well as the characteristics and design considerations related to each of its constituent subsystems. The functions and performance goals of the PODA algorithm were described in terms of DAMA implementation in the WB SATNET. Local area access facilities were described, including the LEXNET and packet voice terminals (PVTs), the Packet Radio net, and circuit-to-packet interfaces, as well as the necessary concentrators and gateways and the associated protocols. Current and planned experimental activity with the wideband system, were described including the integration and performance verification of the major subsystems as well as the near-term goals for experiments combining multiple simultaneous real and emulated conversations over the satellite.

ACKNOWLEDGEMENT

The Wideband Speech Experiment was initiated by Dr. Robert E. Kahn of the DARPA Information Processing Techniques Office. The authors would like to acknowledge his technical contributions as well as his guidance and support through the course of the project.

REFERENCES

- [1] I. M. Jacobs, R. Binder, and E. Hoversten, "General Purpose Packet Satellite Networks," *Proc. IEEE*, Vol. 66, pp. 1448-1467, November 1978.
- [2] T. Bially, A. J. McLaughlin, and C. J. Weinstein, "Voice Communication in Integrated Digital Voice and Data Networks," *IEEE Trans. Comm.*, Vol. COM-28, No. 9, pp. 1478-1490, September 1980.
- [3] C. J. Weinstein, "Statistical Multiplexing of Voice and Data in Broadcast Satellite Networks," *Proc. Fifth International Conference on Computer Comm.*, pp. 811-816, October 1980.
- [4] I. Gitman and H. Frank, "Economic Analysis of Integrated Voice and Data Networks: A Case Study," *Proc. IEEE*, Vol. 66, pp. 1549-1570, November 1978.
- [5] J. W. Forgie, "Voice Conferencing in Packet Networks," *ICC '80 Conference Record*, pp. 21.3.1-21.3.4, June 1980.
- [6] K. Bullington and J. Fraser, "Engineering Aspects of TASI," *Bell System Technical Journal*, Vol. 38, March 1959.
- [7] S. J. Campanella, "Digital Speech Interpolation," *COMSAT Tech. Rev.*, Vol. 6, Spring 1976.
- [8] T. Bially, R. Gold, and S. Seneff, "A Technique for Adaptive Voice Flow Control in Integrated Packet Networks," *IEEE Trans. Comm.*, Vol. COM-28, pp. 325-333, March 1980.
- [9] R. E. Kahn, "The Introduction of Packet Satellite Communications," 1979 NTC Conference Record, pp. 45.1.1-46.1.6, November 1979.
- [10] R. Gold, "Digital Speech Networks," *Proc. IEEE*, Vol. 65, pp. 1636-1658, November 1977.
- [11] S. L. Casner, E. R. Mader, and E. R. Cole, "Some Measurements of ARPANET Packet Voice Transmission," 1978 NTC Conference Record, pp. 12.2.1-12.2.15, November 1978.
- [12] W. W. Chu, et al., "Experimental Results on the Packet Satellite Network," 1979 NTC Conference Record, pp. 45.5.1-45.5.12, November 1979.
- [13] J. W. Forgie, "Speech Transmission in Store-and-Forward Packet Networks," *Proc. NCC* 1975.
- [14] R. F. Sproull and D. Cohen, "High Level Protocols," *Proc. IEEE*, Vol. 66, pp. 1371-1386, November 1978.
- [15] G. C. O'Leary, P. E. Blankenship, J. Tierney, and J. A. Feldman, "A Modular Approach to Packet Voice Terminal Hardware Design," *AFIPS Conference Proceedings (NCC 81)*, Vol. 50, May 1981.
- [16] D. H. Johnson and G. C. O'Leary, "A Local Access Network for Packetized Digital Voice Communication," *IEEE Trans. Comm.*, Vol. COM-29, pp. 679-688, May 1981.
- [17] G. C. O'Leary, "Local Access Facilities for Packet Voice," *Proc. Fifth International Conference Computer Comm.*, pp. 281-286, October 1980.
- [18] P. Spilling and E. Craighill, "Digital Voice Communication in the Packet Radio Network," *ICC '80 Conference Record*, pp. 21.4.1-21.4.7, June 1980.
- [19] M.I.T. Lincoln Laboratory, Semiannual Technical Summary on Information Processing Techniques Program: Packet Speech Systems Technology, ESD-TR-80-71, DDC-AD-A092113/0, March 1980.
- [20] R. Retberg, et al., "Development of a Voice Tunnel System: Design Report," Bolt Beranek and Newman Inc. Report No. 4098, August 1979.
- [21] R. M. Metcalfe and D. R. Boggs, "ETHERNET: Distributed Packet Switching for Local Computer Networks," *Commun. Assoc. Comput. Mach.*, Vol. 19, pp. 395-404, 1976.
- [22] R. E. Kahn, S. A. Gronemeyer, J. Burchfiel, and R. C. Kunzelman, "Advances in Packet Radio Technology," *Proc. IEEE*, Vol. 66, pp. 1468-1496, November 1978.
- [23] J. W. Forgie, "ST - A Proposed Internet Stream Protocol," unpublished memorandum.
- [24] J. B. Postel, "Internetwork Protocol Approaches," *IEEE Trans. Comm.*, Vol. COM-28, pp. 604-611, April 1980.
- [25] M.I.T. Lincoln Laboratory, Annual Report on Network Speech Systems Technology Program, ESD-TR-80-222, DDC-AD-A097998/9, September 1980.
- [26] T. M. Shimabukuro, "The Next Generation COMUS AUTOVON," *IEEE Trans. Comm.*, Vol. COM-28, pp. 1460-1466, September 1980.
- [27] G. J. Coviello and R. H. Levine, "Considerations for a Defense-Switched Network," 1980 NTC Conference Record, November 1980.
- [28] G. Falk, S. Groff, W. M. Milliken, and R. Koosh, "PSAT Technical Report," Bolt Beranek and Newman Inc. Report No. 4469, May 1981.
- [29] I. Jacobs et al., "CPODA - A Demand Assignment Protocol for SATNET," *Proc. Fifth Data Commun. Symposium, ACM and IEEE*, pp. 2-5 to 2-9, September 1977.
- [30] Session 45, 1979 NTC Conference Record, pp. 45.1.1 ff., November 1979.

[51] M.I.T. Lincoln Laboratory, Semiannual
Technical Summary on Information Processing
Techniques Program: Wideband Integrated
Voice/Data Technology, ESD-TR-78-366,
DTIC-AD-A066249/4, September 1978.

This work was sponsored by the Defense Advanced
Research Projects Agency.

The U.S. Government assumes no responsibility
for the information presented.

CERTIFICATE OF SERVICE

I, Andrew A. Lundgren, hereby certify that on October 20, 2006, a true and correct copy of the foregoing document was electronically filed with the Clerk of the Court using CM/ECF, which will send notification that such filing is available for viewing and downloading to the following counsel of record:

Steven J. Balick, Esquire
sbalick@ashby-geddes.com
Ashby & Geddes
222 Delaware Avenue
P.O. Box 1150
Wilmington, DE 19899

I further certify that on October 20, 2006 I caused a copy of the foregoing document to be served by hand delivery on the above-listed counsel of record and on the following non-registered participants in the manner indicated:

BY E-MAIL ON OCTOBER 20, 2006 AND
BY FEDERAL EXPRESS ON OCTOBER 21, 2006

Donald R. Dunner, Esquire
don.dunner@finnegan.com
Finnegan, Henderson, Farabow, Garrett, & Dunner, LLP
901 New York Ave., N.W.
Washington, DC 20001

York M. Faulkner, Esquire
york.faulkner@finnegan.com
Finnegan, Henderson, Farabow, Garrett, & Dunner, LLP
Two Freedom Square
11955 Freedom Drive
Reston, VA 20190

YOUNG CONAWAY STARGATT & TAYLOR, LLP

/s/ Andrew A. Lundgren
John W. Shaw (No. 3362)
jshaw@ycst.com
Andrew A. Lundgren (No. 4429)
alundgren@ycst.com
The Brandywine Building
1000 West Street, 17th Floor
Wilmington, Delaware 19899-0391
(302) 571-6600
Attorneys for Lucent Technologies Inc.